LOADSTAR LETTER

Taking It to Another Level

CMD Prepares to Ship SuperCPUs

LOADSTAR Beta Testing SuperCPU

CMD is in the final stages of testing and tweaking the 20MHZ SuperCPU accelerators. LOADSTAR is lucky enough to be on the short list for beta testing. Check us out next month for details on compatibility and performance.

We'll test the accelerators with LOADSTAR programs and commercial applications and let you know what's

CONTENTS

great and not so great. Considering CMD's amazing success with RAMLink compatibility, we expect

CMD's proposed SuperCPU

smooth sailing. These guys really know what they're doing!

Jeff Fesses Up to Another Bug

By Jeff Jones. Almost from the beginning a few people had major problems with *Starbase 419* by Gus Vakalis. It seemed that anyone who actually beat level two could never coax the program to move to level three.

Well, after a small investigation followed by an attempted cover up, I, Jeff Jones, accept responsibility for adding the bug to Gus' program. It all started when I compressed his forty-level game using

a BASIC pre-incarnation of *RL-EASE*, published on the same issue. When you have a forty-level game and each level takes up four blocks, you start thinking compression.

We'll be republishing *Starbase 419* in archived form. Inside the archive will be the program as supplied by the author. My apologies to Gus for contaminating his program.

Menu Toolbox II Adds High Level Power

By Jeff Jones. On LOADSTAR #144 *Menu Toolbox* will return with BASIC, compiler and machine language level access to *all* commands. So now any programmer can look to the toolbox for prewritten, pretested workhorse routines.

To make programming even less of a hassle, an instant screen command is included. With one SYS, you can set up a screen, background and menu.

LSU-S Creature

By Jeff Jones. It was a biped, intelligent, and enrolled at LSU in Shreveport. It owned a computer -- a Pentium-based machine with a 1.2 gigabyte hard drive. As I spoke with the

creature, it informed me that it owned Photoshop Deluxe, a $500+ program which I purchased legally. It bragged that it was connected with pirates on the Internet and could get hold of anything -- *anything*! Just let it know. It was very excitable.

It then informed me that its computer was currently unusable because it had illegally gathered so much software that his PC had been busy all day crunching its modem's booty. It was hoping that it had room left on its hard drive to actually *use* its computer for some homework.

Gads, I've known pirates, but never one that seemed addicted to a life of crime. If my late-night classroom is transformed into Ruby Ridge by the software police, I'll let you know.

Accelerator Timing Loops

By Jeff Jones. *"If the sound of a tree falling in your program lasts for only a 40th of a second, would anyone hear it?"* The answer is probably no. Well if you've designed a sound to last for a half second on an unaccelerated C-64/128, whether compiled or BASIC, it probably *will* last only a 40th of a second accelerated.

Einstein said it best, *"The rate at which sound moves through time is regulated by crude FOR NEXT loops."* Okay, maybe Einstein never said that, but he understood relativity, and time is very relative in a FOR NEXT loop. Compile or accelerate that FOR NEXT loop and it's suddenly taking much less time to execute. And when your sounds are perfectly timed to exist for a certain number of moments, they become different or even inappropriate sounds when their gate (sounding) is shortened. Cut a bell by four and it becomes a blip. By 20? Maybe you won't even hear a click.

The Folly of Software Delays: Ask the average BASIC programmer to force a half second delay and they'll throw a FOR NEXT loop at you. There is a certain amount of tweaking

involved there to get the desired duration. This isn't the best way to do a delay. Anyone who owns a compiler knows that compiling a program muffs up delay loops and your sounds all sound differently. *TI* is a system variable, and it changes every 60th of a second. The natural way to delay for a half second is to count off 30 jiffies[1]:

```
100 x = ti 110: if ti-x <30
then 110
```

Need a quarter second? Test for less than 15. 3 seconds? <180.

If more games had lines like these instead of arbitrary FOR NEXT loops, they would run perfectly on all systems. The above code will wait for a half second, compiled or uncompiled, accelerated or unaccelerated, or any combination thereof -- with any type of accelerator.

In machine language people tend to increment .X and .Y thousands of times in a loop, tweaking the loop to generate enough of a delay. Unfortunately if you throw that code into an accelerator, it will run too fast. If you throw it into a slower running C-64 emulator, it will run too slowly -- and the delays will be increased accordingly.

In machine language, the thing to do is base your delays on either the hardware clock or the software clock, both of which run at the same speed regardless of the system speed. You could even count RDTIM ($FFDE)[2] and count the changes in the .Y register. Since your interrupt code is called every 60th second, you could set up dedicated timers that set off flags and reset themselves.

---

[1]Jiffy is a measure of time: 1/60th second.

[2].RDTIM is a kernal routine, usually used with machine language programs.

Compiler Timing

By Jeff Jones. I don't expect a huge reduction in FOR NEXT loops because of the previous five paragraphs. These are rules that Fender and I will have to live by since we'll want all of our software to work equally well on all systems.

When your perfectly timed sound is compiled or accelerated, you may run into timing problems. Unfortunately the average delay loop, `ford=1to1000:next`, isn't sped up much by normal compiling. Here's a comparison chart for the code:

```
10 a=ti
20 fori=1to2000:next
30 printti-a
```

Note that I had originally tested a program that counted to 1000, but realized that on a CMD accelerated system with Abacus compiler #2, less than one jiffy would accumulate. Since we need at least one jiffy for testing purposes, I doubled all the numbers. As you can see, the average BASIC delay loop can be accelerated by negligible amounts with normal compilers, but by a factor of four when Abacus' *all-integer* compiler #2 is used. This is because no floating point math is involved. We begin to see a simple way to keep delay loops of the same length whether compiled or uncompiled.

Given that it will always take a normal C-64 124 jiffies to count to 2000, we can use the loop as a test during setup and come up with a factor, *f*, with which we can equalize all subsequent delay loops.

For instance, if we're running a Blitzed program on a CMD 20MHZ accelerated system, the loop will probably take 96/20 or 4.8 jiffies. I'm dividing by 20 since the accelerated program will run 20 times faster. In our test loop, it will probably come out to four jiffies even. Since our standard is 124 jiffies for this loop, we take 124 and divide it by four and come up with a factor of 31. So all delay loops in that

Uncompiled	124
Abacus Compiler #1 P-code	116
Abacus Compiler #1 M-code	112
Blitz!	96
Abacus Compiler #2 P-code	30
Abacus Compiler #2 M-code	30

compiled and accelerated program should be multiplied by a factor of 31. Here is how the factor routine would look:

```
3000 f = ti
3010 fori=1to2000:next
```

```
3020 f = ti - f
3030 if f = ti then f = 1 :
return
3040 f = 124 / f : return
```

Line 3030 exists because the jiffy clock might be off for some reason (you may have previously disabled the STOP key with POKE 788,52). In that case, you can't know how many jiffies have passed so your factor will have to be one. I suggest enabling the STOP key with POKE 788,49 during the test. Now that *f* is defined, we can insert it in all FOR NEXT type delay loops:

```
fori=1to500*f:next
```

will run at roughly the same speed, compiled or uncompiled, accelerated or unaccelerated, or any combination thereof. You can apply this scheme to the sounds supplied under this heading if you choose. I didn't include a factor in the routines since it would make all delays nil unless you first define *f*. That would make me a bit overbearing, wouldn't it? But we do ask that you use such a scheme if you submit a program with sound to LOADSTAR. We do want our programs to work satisfactorily on all systems.

Compiler Integers for Speed

By Jeff Jones. It may seem rather disappointing that compiling a FOR NEXT loop with *Abacus' Basic 64 Compiler* only saves you a couple of jiffies in most circumstances. There's a simple reason for this, and a way to increase your loops many fold.

The easiest way to make your compiled programs run faster is to use compiler #2, plain and simple. Compiler #2 is many times faster than compiler #1. Trouble is you can't always use compiler #2 since it deals only with integers between -32768 and 32767. So you may have to "mix "compilers using compiler directives in REM statements:

```
1000 REM@  01
```

switches to compiler #1 code

```
1070 REM@  02
```

switches to compiler #2 code.

The reasons for the apparent slothfulness of compiler #1 are not numerous. It's all in the floating point math. A straight FOR NEXT delay and PRINT statements aren't sped up much at all by compilers. Your POKEs and computations and things going on *inside* loops will be sped up equally when comparing compilers #1 and #2.

If you can make the loop control variable, *LCV*, an integer variable, you'll greatly enhance the speed of your loops. Trouble is in regular BASIC a line like:

```
FOR I% = 1 TO 1000:NEXT
```

will crash the computer! It's simply not allowed to control a FOR loop with an integer LCV! The compiler allows it though. But if you want your program to run in BASIC as well as compiled, you can't have improper (to BASIC) syntax.

Using compiler #2 makes all variables integer whether you declare them that way or not. That means no variable can be greater than 32767 or less than -32768. If you must use a *real* (floating point) variable inside a compiler #2 program, you can declare variables as real at the start of the program with the following directive:

```
10 REM@R=I,J,K,L
```

This directive makes *I, J, K,* and *L* real variables, even in a compiler #2 program. Using these directives and strategies, you can create the best of both worlds within your compiled program.

The Old ROM Problem

By Jeff Jones. Parts of this article originally appeared on LOADSTAR #85, but I've learned a few things since then. Most important, I've learned how ML programmers can avoid common mistakes that will make their programs impossible to use for people with "old ROM" computers.

Commodore, in its infinite wisdom, has changed the C-64 over time. These changes have meant little to BASIC programmers, but much to ML programmers. The important change is simple: All C-64s DO NOT handle color memory the same way! I'm not going to get into which versions handle color which way. Frankly I can't remember. What's important are *which programming techniques we use to get around the problems.*

Many ML programmers mistakenly believe that if they clear the screen, color memory is filled with the color of the cursor. *We can't afford to be that lazy.* That may be true on your 128 in the 64 mode or any C-64 with JiffyDOS and most true C-64s, but there are some old C-64s that don't fill color memory at all when you clear the screen. There are some that fill color memory with the background color.

Result: You begin poking text and menus and other things to the screen yet *nothing* shows up. The character color is the same as background color. *Then the user calls up LOADSTAR and complains . . .* Oops! Back to programming:

What the old ROM user sees is nothing. Any time you POKE to the screen but don't manually set the color, there are many users who may not be able to see it.

The quick and dirty way to fix this is to set the cursor color to the desired fill color AND set the background to the fill color, then CLR the screen. You would have to poke the screen color back to the normal color when done, which would cause a flash unless you blank the screen.

Talk of such trickery is unsettling my stomach. This is *not* good programming practice! If you're going to poke to the screen, a better way to do it is through the screen and color pointers.

In BASIC screen-poking is usually a no-no because it's slower than a slug on a salt flat. PRINTing in BASIC is about as fast as CHROUT is in ML (at least to the human eye). However in ML, you'll find that

POKEing is a heck of a lot faster than CHROUT because you're optimizing output.

```
LDA #147 ;clear character
JSR CHROUT
JSR OLD'ROM PATCH; explained
later

LDY #119
POKE'BOX LDA #160
STA (209),Y
LDA desired color
STA (243),Y
DEY
BPL POKE'BOX
```

Longer code but faster code. And note that I used 160 to poke to the screen, since we're dealing with screen code, not ASCII.

But what's this (209) and (243)? We never set those pointers -- did we? Yes we did. (209) points to screen memory. Specifically, it points to the beginning of the line the cursor is on (not to the cursor itself.) If the cursor is on the 8th row, the cursor would be at (209),Y with Y being 8. (209) always points to the left HOME of each line.

So when you clear the screen, 209, the low byte of the pointer is zero. 210, the high byte is four.

$$4 * 256 + 0 = 1024$$

Move the cursor down one line? The operating system automatically updates 209 and 210.

$$209 = 40 \quad 210 = 4$$
$$4 * 256 + 40 = 1064$$

That's where row one starts --1064.

Here's where things get hairy! In some old ROM computers the color pointer IS NOT updated by the operating system. Not until a character is printed with CHROUT. When you think about it, it's more efficient, but hackers with late model 64s have no idea that some computers haven't updated this pointer. So if you use the 243 pointer to begin poking to the corresponding color memory, it may work well on YOUR computer, but

some poor guy in Texas with a ten-year-old C-64 watches his computer crash. This is because 243 could be pointing anywhere -- ANYWHERE in memory. Get it? You're poking color right into your program or zero page or your font . . . All because the color pointer hasn't been updated when you moved to the current line.

The way to get around this is to JSR to OLD ROM PATCH whenever you change lines. This is ONLY if you're poking to color memory through the 243 color pointer. If you're printing with CHROUT, you don't need OLD ROM PATCH.

```
OLD'ROM'PATCH LDA 210
SEC
SBC 648
CLC
ADC #$D8
STA 244
LDA #209
STA 243
RTS
```

Okay, okay, don't go grumbling about this taking up too many cycles. This is machine language here! This subroutine is called and returned in approximately 3/100000th sec including the RTS. I don't think it'll make your program crawl!

If you don't call the OLD ROM PATCH after moving between lines, your color pointer may NOT be pointing to the right place. This can cause crashes and theoretically tendonitis and higher taxes.

My famous box routine uses the OLD ROM PATCH, which can be called up to 25 times, depending on the number of lines in your box. Remember, that's only 3/4000th sec added to the drawing time.

Jeff's ML Box Routine

```
        LDX Y1
        LDY #0
```

```
        CLC
        JSR PLOT
        JSR OLD'ROM'PATCH
        INC Y2
        INC X2

....SET'UP LDY X1
. ..DO'BOX LDA CODE
. . STA (209),Y;  POKE REVERSE
SPACE
. . LDA COLOR
. . STA (243),Y;  POKE COLOR
. . INY
. . CPY X2
. ..BCC DO'BOX
.     LDA 214;   WHAT LINE AM I
ON?
.     CMP X2
. ..BEQ QUIT
. . LDA #13
. .  JSR CHROUT; MOVE DOWN A
LINE
. . JSR OLD'ROM'PATCH
....JMP SET'UP
..QUIT RTS

     X1 .BYT 0
     X2 .BYT 0
     Y1 .BYT 0
     Y2 .BYT 0
     CODE .BYT 0
     COLOR .BYT 0
```

This routine will draw a box anywhere on the screen. Note how I used CHROUT to move to the next line, automatically updating (209). This method is very fast. Faster than printing. Printing carriage returns DOES slow down the process by a few thousandths of a second. So if you're in a real hurry to fill the screen, just add #40 to 209 and 210 instead of printing a carriage return. You could also build the OLD ROM PATCH right into the routine so you don't lose cycle time on the JSR and RTS. Seriously though, you shouldn't have to worry about this much efficiency except in raster interrupt work. And this type of code wouldn't be in the raster interrupt routines anyway.

A Machine Language Call To Arms

Machine language programmer, you're a special breed -- even a dying breed. As computers become faster and faster, there will be little need for anyone but operating system authors and compiler authors to know machine language.

If you can learn 6510 assembler then you have a mind meticulous enough to handle a $450 C compiler for your Pentium and secure a future. Don't feel daunted by the thick manuals, and don't let anyone tell you that programming a C-64 is easier than programming an IBM. We have fewer tools right down to the processor level. Frankly, while the C-64 is a very accessible machine, it's a bit harder to program than the IBM. I have whole libraries of standard routines at my disposal when I program the PC. I had to write my own low level routines for the C-64.

While our policy does call for BASIC programs, we respect ML programs. We just hate ML crashes. *Please Don't Be Tricky!* Don't try to give us the smallest, smartest fastest routine. This is a lesson learned well in the PC world. Size doesn't matter. It's the end product *delivering* that counts. We just want it to work for everyone and be noticeably fast. If it must be slower for maximum compatibility, let it be slower. If it must be bigger, let it be bigger.

We also want to be able to edit your ML programs. This may only be to make color or font changes. One way to make editing easier for us is to give us a JMP table that allows us to change the initial setup. You can make this self explanatory by having a boot file that pokes or SYSes parameters to preferences within your program. REM each line so that we know how to change each parameter. That way we don't have to hack your ML at all.

The number one way that I hack ML is to find all the instances of file access that poke hard coded 8s into the current drive before calling the ML. Your ML shouldn't LDX #8 when it uses SETLFS. It should LDX DEVICE which is a variable somewhere in your

program.

The more variables like DEVICE that you create, the more easily configurable your program becomes. You may come up with parameters that no one will ever change, but at least the ones that YOU want to tweak become easy. And if you want to make it user-configurable, you'd have to write it that way anyway.

If you change your mind about a color or position choice, you only have to change the parameters in one place instead of hunting down each instance where you prepped the cursor to print a certain item. When you move to a parameter-driven program, it feels as if you're wasting memory, but you soon realize that while your source code is growing, your object code suffers little. Each new setting only takes up one byte. You could conceivably have a definable setting for every item on your screen and not take up more than 100 bytes of variable space. But what you're left with is a program that we can tweak until it's at its prettiest --without wading through a disassembly listing.

RAM Under ROM

By Jeff Jones. Here's how you can access the RAM under all the ROM, even the IO at $D000-$DFFF. I use these subroutines to knock all ROM images out so that I can get a byte from beneath them, and then on again so that everything continues as normal. Generally I call ROM'OUT right before I need to LDA the data, and call ROM'IN immediately after so that any interrupt driven software suffers the least. These routines don't affect any register.

```
ROM'OUT PHA
LDA #$74
SEI
STA 1
PLA
RTS
```

```
ROM'IN PHA
LDA #$77
STA 1
CLI
PLA
RTS
```

When speed is important, I'll go ahead and write the ROM out/in routines into the larger routine that needs it. I might do whole chunks of work with the ROM out and no interrupts (which is faster), but usually ML is fast enough that I can have a zillion JSRs to ROM'IN and ROM'OUT without anyone noticing.

Remember that ROM must be OUT if you want to write to the RAM between $D000 and $DFFF. The ROM can be IN when writing to the RAM under the BASIC and KERNAL ROMs, but has to be OUT only when reading from them.

Symbolic Assemblers

With assemblers, you can fine tune your code until it shines. I recommend ROCK'S ASSEMBLER, available on the *Compleat Programmer*. See our ad on page 7. I know I don't use it personally, but that's not because I don't like it. It has quite a few features that I wish existed in EBUD. I personally don't use ROCK'S ASSEMBLER because I have become too accustomed to EBUD. I'm just a creature of habit.

Using most any other assembler can cause problems. For instance, MERLIN is widely used, but it's so archaic, cryptic and limited compared to these two assemblers that I can't recommend it. If you're reading this and produce or sell a assembler that you feel is better than EBUD or ROCK'S ASSEMBLER, send it to me and I'll review it. Note that we don't do negative reviews so don't worry about me trashing your product.

There I go digressing again! No one in this age should write code in a monitor or poke ML into memory

from BASIC. If you're going to program ML, do it big and do it right. ML Monitors that allow you to assemble code here and there used to be called assemblers. Now they are considered less than assemblers because they do little to help you edit. They are handy for inspecting and writing little patches, but when you start writing big programs with them, it becomes too much of a problem to fix mistakes. You end up patching in spaghetti-like routines to fix or upgrade, and slowly you forget what's going on in your own code.

Without a symbolic assembler, I'd probably STILL be writing our presenter. There are roughly 70 blocks of text in LOADSTAR BRIEFS. The assembler allowed me to type some, edit it, correct mistakes, assemble whole files of text into the program, and print it.

Machine language programs are no longer merely tiny support modules that serve a larger BASIC program. Large machine language programs are rather commonplace. And these programs are chock full of help and features. I say that if you're going to enhance your BASIC program with ML, or write a 100% ML program, buff it up with color and text. You have more memory available with ML. Use it.